Catastrophic photometric redshift errors: weak lensing survey requirements

Gary Bernstein[1] and Dragan Huterer[2]

[1]Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104
[2]Department of Physics, University of Michigan, 450 Church St, Ann Arbor, MI 48109-1040

We study the sensitivity of weak lensing surveys to the effects of *catastrophic* redshift errors — cases where the true redshift is misestimated by a significant amount. To compute the biases in cosmological parameters, we adopt an efficient linearized analysis where the redshift errors are directly related to shifts in the weak lensing convergence power spectra. We estimate the number $N_{\rm spec}$ of unbiased spectroscopic redshifts needed to determine the catastrophic error rate well enough that biases in cosmological parameters are below statistical errors of weak lensing tomography. While the straightforward estimate of $N_{\rm spec}$ is $\sim 10^6$, we find that using only the photometric redshifts with $z \lesssim 2.5$ leads to a drastic reduction in $N_{\rm spec}$ to $\sim 30,000$ while negligibly increasing statistical errors in dark energy parameters. Therefore, the size of spectroscopic survey needed to control catastrophic errors is similar to that previously deemed necessary to constrain the core of the $z_s - z_p$ distribution. We also study the efficacy of the recent proposal to measure redshift errors by cross-correlation between the photo-z and spectroscopic samples. We find that this method requires $\sim 10\%$ *a priori* knowledge of the bias and stochasticity of the outlier population, and is also easily confounded by lensing magnification bias. The cross-correlation method is therefore unlikely to supplant the need for a complete spectroscopic redshift survey of the source population.

I. INTRODUCTION

Weak gravitational lensing is a very promising cosmological probe that has potential to accurately map the distribution of dark matter and measure the properties of dark energy and the neutrino masses (for reviews, see [1–4]). It is well understood, however, that systematic errors may stand in the way of weak lensing reaching its full potential — that is, achieve the *statistical* errors predicted for future ground and space based surveys such as the Dark Energy Survey (DES), Large Synoptic Survey Telescope (LSST), and the Joint Dark Energy Mission (JDEM). Controlling the systematic errors is a primary concern in these and other surveys so that a variety of dark energy tests (recently proposed and reviewed by the JDEM Figure of Merit Science Working Group [5]) can be performed to a desired high accuracy.

Several important sources of systematic errors in weak lensing surveys have already been studied. Chief among them is the redshift accuracy—approximate, photometric redshifts are necessary because it is infeasible to obtain optical spectroscopic redshifts for the huge number ($\sim 10^8$-10^9) of galaxies that future surveys will utilize as lensing sources. It is therefore imperative to ensure that statistical errors and systematic biases in the relation between photometric and spectroscopic redshifts (recently studied in depth with real data [6–10]) do not lead to appreciable biases in cosmological parameters.

The relation between the photometric and spectroscopic redshift has been previously modeled as a Gaussian with redshift-dependent bias and scatter. It is found that both the bias and scatter (that is, quantities $\langle z_p - z_s\rangle$ and $\langle (z_p - z_s)^2\rangle^{1/2}$ in each bin of $\Delta z = 0.1$), need to be controlled to about 0.003 or better in order to lead to less than $\sim 50\%$ degradation in cosmological parameter accuracies [11–13]. These constraints are a bit more stringent in the most general case when the redshift error cannot be described as a Gaussian [14]. These requirements imply that $N_{\rm spec} \lesssim 10^5$ spectra are required in order to calibrate the photometric redshifts to the desired accuracy. Gathering such relatively large spectroscopic sample will be a challenge, setting a limit to the useful depth of weak lensing surveys. [While the redshift errors have been well studied, other systematics are also important, especially theoretical errors in modeling clustering of galaxies at large and small scales, intrinsic shape alignments, and various systematic biases that take place during observations [15–50].]

All of the aforementioned photo-z requirement studies (e.g. [11, 12]), however, have modeled the errors as a perturbation around $z_s - z_p$ relation. While this perturbation was allowed to be large and to have a nonzero scatter and even skewness (e.g. [11]), it did not subsume a general, multimodal error in redshift.

In this paper we would like to remedy this omission by estimating the effect of *catastrophic* redshift errors. Catastrophic errors are loosely defined as cases when the photometric redshift is grossly misestimated, i.e. when $|z_p - z_s| \sim O(1)$, and are represented by arbitrary "islands" in the $z_p - z_s$ plane. We develop a formalism that treats these islands as small "leakages" (or "contaminations") and directly estimates their effect on bias in cosmological parameters. We then invert the problem by estimating how many spectroscopic redshifts are required to control catastrophic errors at a level that makes them harmless for cosmology.

The paper is organized as follows. In §II we derive the relevant equations for the bias in cosmological parameters induced by misestimated catastrophic redshift errors in a tomographic weak lensing survey. In §III we apply these methods to a canonical ambitious weak-lensing cosmology project. In §IV we ask: how large must a *complete* spectroscopic redshift survey be in order that the catastrophic photo-z error rates be measured sufficiently well that remnant cosmological biases are well below the statistical uncertainties? Newman [51] has suggested an alternative mode of measuring the photo-z error distribution, namely the angular cross-correlation of the photometric galaxy sample nominally at z_p with a spectroscopic sample at z_s; in §V we investigate whether systematic errors in the photo-z outlier rates derived from this cross-correlation technique will be small enough to render cosmological biases insignificant. The final section discusses the scaling of these results with critical survey parameters, the ramifications for survey design, and areas of potential future investigation.

II. FORMALISM

In this section we establish the formalism that takes us from "islands" in the $z_s - z_p$ plane to biases in cosmological parameters. First, however, we define the basic observable quantity, the convergence power spectrum, and its corresponding Fisher information matrix.

A. Basic observables and the Fisher matrix

The convergence power spectrum of weak lensing at a fixed multipole ℓ and for the ith and jth tomographic bin is given by

$$\mathcal{P}_{ij}^{\kappa}(\ell) = \frac{\ell^3}{2\pi^2}\int_0^{\infty} dz\, \frac{W_i(z)\, W_j(z)}{r(z)^2\, H(z)}\, \mathcal{P}_{\rm mat}\left(\frac{\ell}{r(z)}, z\right), \tag{1}$$

where $r(z)$ is the comoving distance, $H(z)$ is the Hubble parameter, and $\mathcal{P}_{\rm mat}(k, z)$ is the matter power spectrum. The weights W_i are given, for a flat Universe, by $W_i(\chi) = \frac{3}{2}\,\Omega_M\, H_0^2\, g_i(\chi)\,(1+z)$ where $g_i(\chi) = \chi\int_{\chi}^{\infty} d\chi_s n_i(\chi_s)(\chi_s - \chi)/\chi_s$, χ is the comoving distance and n_i is the comoving density of galaxies if χ_s falls in the distance range bounded by the ith redshift bin and zero otherwise. We employ the redshift distribution of galaxies of the form $n(z) \propto z^2\exp(-z/z_0)$ that peaks at $2z_0 \simeq 0.9$.

The observed convergence power spectrum is

$$C_{ij}^{\kappa}(\ell) = \mathcal{P}_{ij}^{\kappa}(\ell) + \delta_{ij}\frac{\langle\gamma_{\rm int}^2\rangle}{\bar{n}_i}, \tag{2}$$

where $\langle\gamma_{\rm int}^2\rangle^{1/2}$ is the rms intrinsic shear in each component which we assume to be equal to 0.24, and $\bar{n}_i$ is the average number of galaxies in the ith redshift bin per steradian. The cosmological constraints can then be computed from the Fisher matrix

$$F_{ij} = \sum_{\ell}\frac{\partial\mathbf{C}}{\partial p_i}\,\mathbf{Cov}^{-1}\,\frac{\partial\mathbf{C}}{\partial p_j}, \tag{3}$$

where $\mathbf{Cov}^{-1}$ is the inverse of the covariance matrix between the observed power spectra. For a Gaussian convergence field, its elements are given by

$$\mathrm{Cov}\left[C_{ij}^{\kappa}(\ell), C_{kl}^{\kappa}(\ell)\right] = \frac{\delta_{\ell\ell'}}{(2\ell+1)\, f_{\rm sky}\,\Delta\ell}\,\left[C_{ik}^{\kappa}(\ell)C_{jl}^{\kappa}(\ell) + C_{il}^{\kappa}(\ell)C_{jk}^{\kappa}(\ell)\right]. \tag{4}$$

where $f_{\rm sky}$ is fraction of the sky observed and $\Delta\ell$ is the binning of the convergence power spectra in multipole.

Our fiducial SNAP survey described below, without any theoretical systematics, determines w_0 and w_a to accuracies of $\sigma(w_0) = 0.089$ and $\sigma(w_a) = 0.31$ (corresponding to the pivot value determined to $\sigma(w_p) = 0.027$).

B. Biases in the Gaussian limit

Consider the general problem of constraining a vector of cosmological parameters $P = \{p_i\}$ based on an observed data vector $D = \{D_\alpha\}$. If the observable quantities D_α are distributed as Gaussians with covariance matrix C, then the first-order formula for bias in the ith parameter, Δp_i, induced by a bias ΔD in the data is (e.g. [52, 53])

$$\Delta p_i = \sum_j (F^{-1})_{ij} \sum_{\alpha\beta} \frac{\partial \bar{D}_\alpha}{\partial p_j} (C^{-1})_{\alpha\beta}\, \Delta D_\beta. \tag{5}$$

Here F is the Fisher matrix for the cosmological parameters, and is defined as the second derivative of log likelihood ($\mathcal{L} \equiv -\ln L$) with respect to the parameters. The bias above can be more concisely expressed as

$$\Delta P = F^{-1} Q\, \Delta D \equiv F^{-1} V, \tag{6}$$

where we have defined the matrix Q and vector V as

$$Q_{ij} \;\equiv\; \sum_k \frac{\partial \bar{D}_k}{\partial p_i} (C^{-1})_{kj} \tag{7}$$

$$V \;\equiv\; Q\, \Delta D. \tag{8}$$

The induced parameter bias is considered unimportant if it is small compared to the expected statistical variation in the cosmological parameters. In the case where the likelihood in the parameter space is Gaussian, the likelihood of the bias ΔP being exceeded by a statistical fluctuation is determined by

$$\Delta\chi^2 = \Delta P^T F \Delta P = V^T F^{-1} V \tag{9}$$

In the Appendices we prove two useful theorems about $\Delta\chi^2$:

1. The bias significance $\Delta\chi^2$ always *decreases* or stays fixed when we augment the likelihood with (unbiased) prior information, *e.g.* data from a non-lensing technique;

2. $\Delta\chi^2$ always *decreases* or stays fixed when we marginalize over one or more dimensions of the parameter space. In the Gaussian limit, the bias Δp_i is unaffected by marginalization over other parameters.

We will use these results later to argue that our requirements on the control of catastrophic redshift errors are conservative, in the sense that adding other cosmological data or considering individual cosmological parameter biases will only weaken the requirements.

C. The case of catastrophic photo-z errors

For weak-lensing tomography, the data elements D_α are the convergence (or shear) cross-power spectrum elements $C^{\kappa}_{\alpha\beta}(\ell)$ between photo-z bins α and β at multipole ℓ. Let us examine how these will be biased by photo-z outliers. [The data covariance matrix C of §II B is the matrix Cov of Eq. (4).]

We assume a survey with the (true) distribution of source galaxies in redshift $n_S(z)$, divided into some number N_b of bins in redshift. Let us define the following terms

- *Leakage:* fraction of objects from a given spectroscopic bin that are placed into an incorrect (non-corresponding) photometric bin.

- *Contamination:* fraction of galaxies in a given photometric bin that come from a non-corresponding spectroscopic bin.

One could estimate either of these quantities — after all, when specified for each bin, they contain the same information. Let *leakage* fraction l_{ST} of galaxies in some spectroscopic-redshift bin n_S (the "source" of leakage) end up in some photo-z bin n_T (the "target" of leakage), so that l_{ST} is the fractional perturbation in the source bin. Note that, since bins S and T may not have the same number of galaxies, the fractional perturbation in the target bin is not the same. The *contamination* of the target bin T, c_{ST}, is related to the source bin leakage via

$$c_{ST} = \frac{N_S}{N_T}\, l_{ST} \tag{10}$$

where N_S and N_T are the absolute galaxy numbers in the source and target bin respectively.

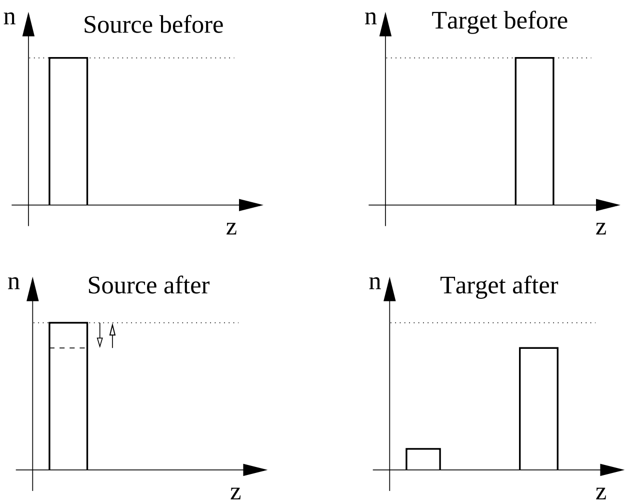

FIG. 1: Explanation of how the leakage and contamination operate. In this figure, we assume for simplicity that the number of galaxies in the source and target bin is the same, so that $l_{ST} = c_{ST}$. Because the redshift distribution $n(z)$ is normalized to unit integral in each bin, the source bin's redshift distribution $n_S(z)$ does not change; see the bottom left panel. The target bin's redshift distribution, $n_T(z)$, does change however, as illustrated in the bottom right panel.

The redshift distribution of galaxies (normalized to unity) in the source bin, n_S, does not change, since a fraction of galaxies is lost — but the redshift distribution is normalized to unit integral; see Fig. 1. Conversely, things are perturbed in the target bin, since it now contains two populations of galaxies, the original one with fraction $1 - c_{ST}$, and the contamination at incorrect (source bin) redshift with fraction c_{ST}; again this is clearly shown in Fig. 1. Therefore

$$n_S \;\to\; n_S \tag{11}$$

$$n_T \;\to\; (1 - c_{ST})\, n_T + c_{ST}\, n_S \tag{12}$$

and only the *target* bin is affected (i.e. biased) by photo-z catastrophic errors.

The effect on the cross power spectra is now easy to write down. Clearly, only the (α, β) cross-spectra where one of the bins is the target bin — $\alpha = T$ or $\beta = T$ — will be affected

$$\begin{aligned} C_{TT} &\;\to\; (1 - c_{ST})^2 C_{TT} + 2c_{ST}(1 - c_{ST})C_{ST} + c_{ST}^2 C_{SS} \\ C_{\alpha T} &\;\to\; (1 - c_{ST})C_{\alpha T} + c_{ST}\, C_{\alpha S} \qquad (\alpha \neq T) \\ C_{\alpha\beta} &\;\to\; C_{\alpha\beta} \qquad\qquad\qquad\qquad\quad \text{(otherwise)} \end{aligned} \tag{13}$$

We have checked that ignoring the quadratic terms in $c_{\alpha\beta}$ leads to no observable effects to the results (for $c_{ST} = 0.001$ contamination). The biases can now be computed as the right hand side minus the left hand side in the formulae above. We replace the single index α for data elements in Eq. (5) with the triplet $\ell\alpha\beta$ so that $D_{\ell\alpha\beta} \equiv C^{\kappa}_{\alpha\beta}(\ell)$, and we reserve the symbol C for the covariance of the data elements. The bias in data induced by the catastrophic errors is

$$\Delta D_{\ell\alpha\beta} = \sum_{\mu\nu} c_{\mu\nu} \left[\delta_{\alpha\nu}(D_{\ell\mu\beta} - D_{\ell\alpha\beta}) + \delta_{\beta\nu}(D_{\ell\mu\alpha} - D_{\ell\alpha\beta})\right]. \tag{14}$$

If we make the further assumption that the convergence is a Gaussian random field, then we have

$$C_{\ell\alpha\beta,\ell'\gamma\delta} \;=\; \delta_{\ell\ell'} \left[D_{\ell\alpha\gamma} D_{\ell\beta\delta} + D_{\ell\alpha\delta} D_{\ell\beta\gamma}\right] \tag{15}$$

$$\Rightarrow \quad (C^{-1})_{\ell\alpha\beta,\ell'\gamma\delta} \;=\; \frac{\delta_{\ell\ell'}}{2} (D_\ell^{-1})_{\alpha\gamma} (D_\ell^{-1})_{\beta\delta}. \tag{16}$$

Eq. (5) simplifies considerably when we invoke Eqs. (14) and (16):

$$\Delta p_i = \sum_{j,\mu\nu} (F^{-1})_{ij} M_{j,\mu\nu} c_{\mu\nu}, \tag{17}$$

$$M_{j,\mu\nu} \equiv \sum_\ell \left[(E_i^\ell)_{\mu\nu} - (E_i^\ell)_{\nu\nu} \right], \tag{18}$$

$$E_i^\ell \equiv \frac{\partial D_\ell}{\partial p_i} D_\ell^{-1}. \tag{19}$$

As a reminder, the Fisher matrix in the case of a zero-mean Gaussian variable is [54]

$$F_{ij} = \frac{1}{2} \sum_\ell \mathrm{Tr}(E_i^\ell E_j^\ell). \tag{20}$$

In a cosmological application we will marginalize over all parameters except a subset of interest A. In the Fisher approximation bias is simply projected onto the A subset: $\Delta p_A = P_A \Delta p$, where P_A is the projection matrix (see Appendix B). If F_A is the marginalized Fisher matrix, then the $\Delta\chi^2$ of the bias after marginalization is

$$\Delta\chi^2 = \sum_{\mu\nu} \sum_{\mu'\nu'} c_{\mu\nu} c_{\mu'\nu'} \sum_{ij} M_{i,\mu\nu} \left(F^{-1} P_A^T F_A P_A F^{-1} \right)_{ij} M_{j,\mu'\nu'}. \tag{21}$$

III. APPLICATION TO CANONICAL SURVEYS

For first study we examine a weak lensing survey similar to that proposed for SNAP [55], but with the source-galaxy selection restricted to incur minimal catastrophic error rate. Evaluation of other potential surveys could be performed following the same model.

We take the eight-parameter cosmological model considered by the Dark Energy Task Force (DETF; [56]): dark energy physical density $\Omega_{\rm DE} h^2$, and equation of state parameters w_0 and w_a; normalization of the primordial power spectrum A and spectral index n; and matter, baryon, and curvature physical densities $\Omega_M h^2, \Omega_B h^2$, and $\Omega_k h^2$. The fiducial values of these parameters are taken from the 5-year WMAP data [57]. We will assume a Planck CMB prior as specified by the DETF report. Recall that application of further priors can only weaken the requirements on photo-z outliers (see Appendix A).

We assume shear tomography with 20 bins linearly spaced over $0 < z < 4$ with $\Delta z = 0.2$; we have checked that the results are stable with Δz. The redshift distribution and fiducial $c_{\alpha\beta}$ are taken from a simulation of the photo-z performance of SNAP as described in [58]. The procedure is to (1) create a simulated catalog of galaxies; (2) calculate their noise-free apparent magnitudes in the SNAP photometric bands spanning the visible and NIR to 1.6 μm; (3) add the anticipated observational noise to each magnitude; (4) determine a best-fit galaxy type and redshift with the template-fitting program *Le Phare*[1]; (5) examine the 95% confidence region $z_l < z_p < z_h$ determined by *Le Phare* and retain only those galaxies satisfying [59]

$$D95 \equiv \ln \left(\frac{1 + z_h}{1 + z_l} \right) \leq 0.15. \tag{22}$$

This strict cut results in a catalog of ≈ 70 galaxies per arcmin^2, with a median redshift of $z_m \approx 1.2$. The WL survey is assumed to cover $f_{\rm sky} = 0.1$ of the full sky, with shape noise of $\sigma_\gamma = 0.24$ per galaxy. We consider only shear tomography at the 2-point level, as this will likely maximize the bias imparted by redshift errors. We also ignore systematic errors other than redshift outliers, which will likely maximize the statistical significance of the outlier bias.

We will henceforth in this paper define a redshift outlier to satisfy

$$\left| \ln \frac{1 + z_p}{1 + z_s} \right| > 0.2 \qquad \text{(catastrophic outlier definition)}. \tag{23}$$

In the simulated photo-z catalog, 2.2% of the source galaxies are outliers by this criterion. In our analyses below we will only consider biases from photo-z errors meeting this outlier criterion, i.e. we assume the "core" of the photo-z distribution is well determined.

---

[1] www.oamp.fr/people/arnouts/LE PHARE.html

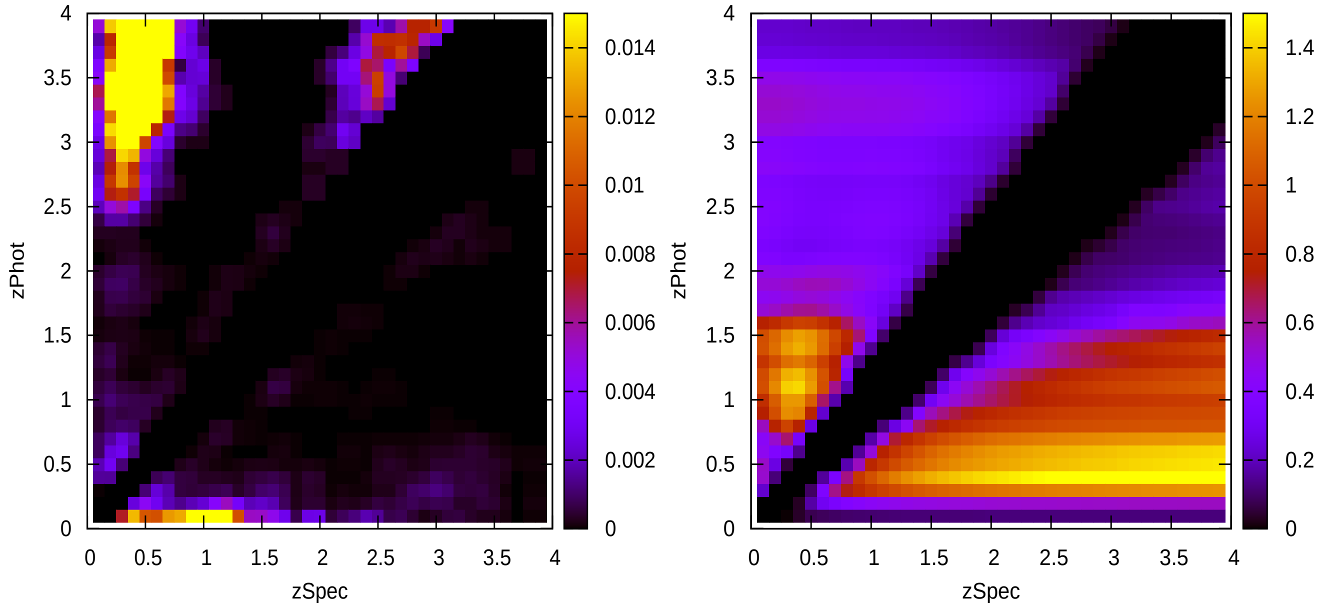

FIG. 2: *Left:* The contamination rate $c_{sp}/\Delta z_s$ of the photo-z bin per unit redshift in spectro-z is plotted for our example case. Note that the contamination is highest at $z_p > 2.5$ and $z_P < 0.2$, where there are relatively few source galaxies and hence a small number of outliers can become a large fractional contamination. *Right:* The quantity h_{sp} which specifies the significance of the $w_0 - w_a$ bias caused by a contamination rate of $0.001\,\delta z_p$ across a range δz_p of photo-z. This plot indicates that the outlier contamination rates must be known to 1–3 parts per thousand over all photometric redshift bins, most sensitively at $0.3 < z_p < 1.5$.

Figure 2 illustrates the canonical model and the sensitivity to redshift outliers in this model. The left-hand panel shows the quantity $c_{sp}/\Delta z$ vs z_s and z_p. [We scale the contamination by Δz to produce a quantity that is independent of the choice of bin size Δz.] The highest contaminations are in two "islands": one at $z_p > 2.5$, $z_s < 0.6$ is probably due to confusion between high-z Lyman breaks and low-z 400-nm breaks. Because true $z > 2.5$ galaxies are relatively rare, a small leakage rate from $z_s \sim 0.5$ can produce a high contamination fraction. The *Le Phare* code run for this simulation does not incorporate a magnitude prior for the photo-z; doing so might reduce the size of this island.

A second high-contamination region is $z_s \approx 1$, $z_p < 0.2$. Again the contamination rate is high because the target-bin density is much lower than the source-bin density.

The right-hand panel in Figure 2 shows $\Delta\chi^2$ evaluated using Eq. (21) for this case of catastrophic errors. We calculate the significance $\Delta\chi^2_{2d}$ of the bias after marginalization of the cosmology onto the $w_0 - w_a$ plane. We simplify by considering the bias arising from contamination in a single bin. This is

$$\Delta\chi^2_{2d} = (h_{\mu\nu}\Delta z)^2 c^2_{\mu\nu}, \tag{24}$$

$$h^2_{\mu\nu} \equiv \sum_{ij} M_{i,\mu\nu} \left(F^{-1}P_A^T F_A P_A F^{-1}\right)_{ij} M_{j,\mu\nu}/(\Delta z)^2. \tag{25}$$

Again the inclusion of the Δz factor defines a $h_{\mu\nu}$ that is invariant under rebinning. The interpretation of $h_{\mu\nu}$ is as follows: if there is an "island" of outliers that spans a range δz_p of photo-z bins, and contains a fraction $\bar{c}$ of the galaxies in these photo-z bins, then the 2d significance of the resultant bias will be

$$\sqrt{\Delta\chi^2} \approx h_{\mu\nu}\delta z_p \bar{c}. \tag{26}$$

Figure 2 has been scaled by 1000, so that it indicates the bias significance of a contamination rate $\bar{c} = 0.001/\delta z_p$. We desire $\Delta\chi^2_{2d} \ll 2.3$ to keep the bias well within the 68% confidence contour. The most severe constraint on $\bar{c}$ would be to take the peak value $h_{\mu\nu} \approx 1300$ and a very wide island, $\delta z_p \approx 0.5$, in which case the criterion for small bias becomes

$$\bar{c} < 1/(h_{\mu\nu}\delta z_p) \approx 1/(1300 * 0.5) = 0.0015. \tag{27}$$

The contamination rate into any island of outliers must be known to 0.0015 or better to avoid significant cosmological bias. This conclusion is independent of the nominal outlier rate. The tolerance on outlier rate will scale with sky coverage as $f_{\rm sky}^{-1/2}$.

IV. CONSTRAINT VIA SPECTROSCOPIC SAMPLING

The most obvious way to determine the contamination rate $c_{\alpha\beta}$ is to conduct a *complete* spectroscopic redshift survey of galaxies in photo-z bin β. It is of course essential that the spectra be of sufficient quality to determine redshifts even for the outliers in the sample.

Let us now estimate the total number of spectra $N_{\rm spec}$ required in order to keep the total bias below some desired threshold. We will assume that each redshift drawn from the spectroscopic survey is statistically independent. In this case the distribution of $N_{\alpha\beta}$, the number of galaxies in photo-z bin β that have spectro-z in bin α, will be described by a multinomial distribution. When the outlier rates are small, the number of spectra in each outlier bin tend toward independent Poisson distributions.

We would like to relate the contamination uncertainties δc_{ij} to the required number of galaxy spectra. Let N_β be the number of spectra drawn from the photometric redshift bin β so that $N_{\rm spec} = \sum_\beta N_\beta$. In this case $\langle N_{\alpha\beta}\rangle = c_{\alpha\beta} N_\beta$ and the variance of the contamination estimate is

$$\delta c_{\alpha\beta}^2 = \frac{(\delta N_{\alpha\beta})^2}{\langle N_\beta\rangle^2} = \frac{c_{\alpha\beta}}{N_\beta} \tag{28}$$

Since the Poisson errors between different outlier bins are uncorrelated, the expected bias significance becomes

$$\langle \Delta\chi^2\rangle = \sum_{\alpha\beta} (h_{\alpha\beta}\Delta z)^2 \langle \delta c_{\alpha\beta}^2\rangle = \sum_{\alpha\beta} (h_{\alpha\beta}\Delta z)^2 c_{\alpha\beta}/N_\beta. \tag{29}$$

We would like to quote a total number of spectroscopic redshifts rather than the number per photo-z slice (N_β above) in order to make our findings more transparent. We consider two cases: first, a slitless or untargeted spectroscopic survey will obtain redshifts in proportion to the number density n_β of source galaxies in each redshift bin: $N_\beta = N_{\rm spec} n_\beta/n$. Then we will consider a targeted survey, in which the number N_β can be chosen bin-by-bin to produce the minimal bias for given total $N_{\rm spec}$.

A. Untargeted spectroscopic survey

In the untargeted case, $N_\beta = N_{\rm spec} n_\beta/n$ and the condition $\Delta\chi^2 \ll 1$ becomes (from Eq. (29))

$$N_{\rm spec} \gg \sum_{\alpha\beta} (\Delta z)^2\, h_{\alpha\beta}^2 c_{\alpha\beta} n/n_\beta. \tag{30}$$

Figure 3 plots the summand of this expression in the $z_s - z_p$ plane. The required $N_{\rm spec}$ is hence the sum over values in this plane (note that we omit the bins near the diagonal that do not meet the "outlier" definition). We find that $\Delta\chi^2_{2d} \lesssim 1$ requires $N_{\rm spec} \gtrsim 8\times 10^5$, and in the full 8-D parameter space (that is, considering $\Delta\chi^2$ for the 8-dimensional parameter ellipsoid), we need $N_{\rm spec} \gtrsim 2\times 10^6$.

These requirements are daunting, potentially larger than the $N_{\rm spec}$ that are needed to constrain the core of the photo-z distribution as determined by Ma and Bernstein [14]. Note however that the requirement is strongly driven by the region $z_s > 2.5$, $z_p < 0.8$. This is because contamination rates are high here in the nominal photo-z distribution (Figure 2), and these bins are sparsely populated (small n_p), meaning that many spectra must be taken in order to accumulate a strong enough constraint on these contamination coefficients.

This suggests a strategy of omitting the $z_p > 2.5$ galaxies from the tomography analysis entirely. Omitting $z_p > 2.5$ from the sum (30) produces a much relaxed requirement: for $\Delta\chi^2_{2d} \lesssim 1$, we need $N_{\rm spec} \gtrsim 2.8\times 10^4$, a 30-fold reduction. This strategy would eliminate $\approx 8\%$ of the source galaxies in the SNAP model, and reduce the dark-energy constraint power by 18%, as mesured by the DETF figure of merit. This would be an acceptable strategy to reduce the outlier bias if one were unable to eliminate the high-z_p island of outliers by refinements to the photo-z methodology.

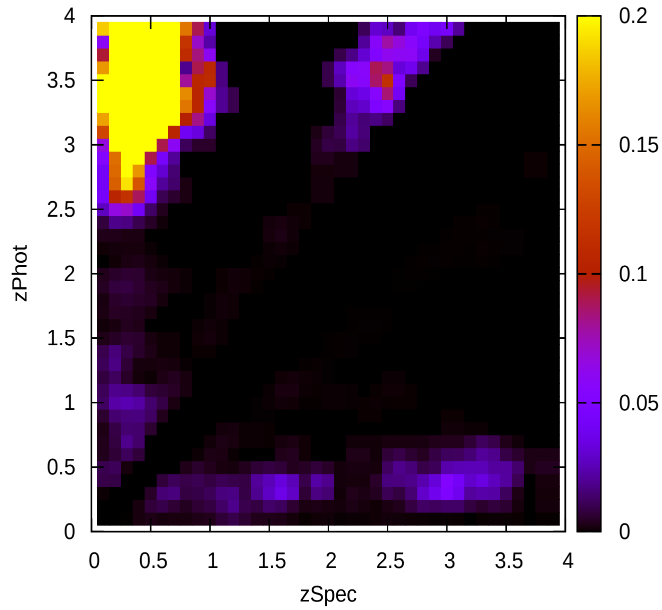

FIG. 3: Number of calibration spectra required to attain an outlier bias significance of $\Delta\chi^2_{2d} = 1$ is given by the integral of this plot over the $z_s - z_p$ space. The quantity plotted is, from Eq. (30), $h^2_{sp} c_{sp} n/n_p$, where n_p is the source density in the photo-z bin and n is the total source density. The plot is in units of 10^6 galaxies. Note that the required number of calibration spectra is strongly driven by the need to constrain the outliers with photo-z $z_p > 2.5$ but true redshifts $z_s < 0.8$.

B. Targeted spectroscopic survey

If we wish to minimize $\Delta\chi^2$ in Eq. (29) for a given total $N_{\rm spec}$, a simple optimization yields

$$N_\beta \propto \sqrt{\sum_\alpha h^2_{\alpha\beta} c_{\alpha\beta}}, \tag{31}$$

$$N_{\rm spec}\Delta\chi^2 = \left(\sum_\beta \Delta z \sqrt{\sum_\alpha h^2_{\alpha\beta} c_{\alpha\beta}}\right)^2 \tag{32}$$

Optimized targeting reduces the requirement for $\Delta\chi^2_{2d} \lesssim 1$ to be $N_{\rm spec} \gtrsim 1.2 \times 10^5$. Eliminating the $z_p > 2.5$ sources reduces the requirement sixfold, $N_{\rm spec} \gtrsim 2.0 \times 10^4$.

Note that the targeted redshift requires 7 times fewer calibration redshifts than the untargeted survey, if we are using the full source redshift range, but only 1.4 times smaller $N_{\rm spec}$ if we restrict $z_p < 2.5$ in the lensing analysis.

C. Scaling and Robustness

The required $N_{\rm spec}$ to reduce outlier-rate biases to insignificance scales with the sky coverage and the mean outlier rate as

$$N_{\rm spec} \propto f_{\rm sky}\bar{c} \tag{33}$$

when most of the information is coming from shear tomography, and the depth of the survey is held fixed. We have verified that $N_{\rm spec}$ varies little as the number of tomography bins grows ($\Delta z < 0.2$) and the information content of the tomography saturates. The two bias theorems imply that the required $N_{\rm spec}$ will drop if we add additional unbiased prior information, or if we marginalize down to a single dark-energy parameter.

More precisely, we find that the ratio $N_{\rm spec} \times (\sigma(w_p) \times \sigma(w_a)) \equiv N_{\rm spec}/{\rm FoM}$ (featuring the well known "figure of merit" [53, 56]) is roughly the same with several alternative survey specifications we consider (and is exactly the same if only the sky coverage $f_{\rm sky}$ is varied).

We have used two independent codes to verify the robustness of results to the myriad of assumptions made and check for the presence of unwanted numerical artifacts. The two codes agree to roughly a factor of two in $N_{\rm spec}$, which is satisfactory given the differences between the implementations, *e.g.* whether curvature and/or neutrino masses are free to vary, and whether the fiducial redshift distribution is smoothed over the cosmic variance in the simulated catalog.

D. Correlated outlier errors and incompleteness

So far we have considered the case where bin-to-bin fluctuations in contamination $\delta c_{\alpha\beta}$ are uncorrelated. Contamination-rate errors $\delta c_{\alpha\beta}$ that are *correlated* from bin to bin might arise if the spectroscopic survey systematically misses outliers in certain redshifts islands (or if the spectroscopy is not done at all!). We can set a specification on the maximum allowable systematic contamination rate error $\delta\bar{c}$ in an island of photo-z width δz_p using Eq. (26). Our previous results for the canonical survey show that the contamination rate in the island should be known to $\delta\bar{c} \lesssim 0.0015$. In other words a spectroscopic-redshift failure rate of only 0.15% in some range of z_p can cause a significant cosmology bias if all of these missed redshift are outliers in a particular island. A 99.9% success rate has rarely if ever been achieved in a spectroscopic redshift survey.

V. CONSTRAINING OUTLIER RATES USING GALAXY CORRELATIONS

The above requirements on $N_{\rm spec}$ and completeness may be too expensive to accomplish, particularly for fainter galaxies. We now examine the possibility that one could make use of a spectroscopic galaxy sample that does *not* fairly sample the photo-z galaxies [51]. The idea is to cross-correlate a photo-z sample at nominal redshift z_p with a spectroscopic sample known to be confined to a distinct bin z_s. The amplitude of this cross-correlation will tell us something about the contamination rate c_{sp}, since there is no intrinsic correlation between the galaxy densities at the two disparate redshifts.

Newman [51] calculates the errors in this estimate that would be induced by shot noise in the sample (for a somewhat related work, see [60]). Here we assume that statistical errors will be negligible and attend to two systematic errors that will arise.

A. Outlier bias and correlation coefficients

First, we define g_s, g_p to be the fluctuations in sky density of the projected distributions of the spectroscopic-survey galaxies in bin s and the photometric-redshift galaxies in bin p. We set m_s to be the fractional fluctuations in the projected mass density in redshift bin s. The bias is defined by $\langle g_s^2\rangle = b_s^2\langle m_s^2\rangle$. The amplitude of the measured cross-correlation between the angular distributions g_s and g_p of the spectroscopic and photometric samples can be written as

$$\langle g_s g_p\rangle = c_{sp} b_s b_{sp} r_{sp}\langle m_s^2\rangle = c_{sp}\frac{b_{sp}r_{sp}}{b_s}\langle g_s^2\rangle, \tag{34}$$

where b_{sp} is the bias of the *outlier population*, and r_{sp} is the correlation coefficient between the density fluctuations of the spectroscopic and the outlier populations. The outlier population are those minority of sources in photo-z bin p that have spectroscopic redshifts in bin s.

In a large survey, shot noise in $\langle g_s g_p\rangle$ and in $\langle g_s^2\rangle$ might become small, and b_s could be determined to good accuracy through study of the redshift-space power spectrum of the spectroscopic targets. It will be difficult, however, to discern b_{sp} because the angular correlation signal of the outlier population is overwhelmed by the correlations of the galaxies in the "core" of the photo-z error distribution (i.e. photo-zs which have *not* been catastrophically misestimated). The correlation coefficient r_{sp} also has no alternative observable signature we have identified. Hence there will be a systematic-error floor on δc_{sp} arising from the finite *a priori* knowledge of the product $b_{sp}r_{sp}$.

B. Lensing Magnification

The second complication to the cross-correlation method is that gravitational lensing magnification bias will induce a correlation between the spectroscopic and photometric samples even if there is *no* contamination. Let us assume that the spectroscopic sample is in the foreground of the photo-z "core"; a similar analysis can be done when using cross-correlation to search for contamination by background galaxies. There are two types of magnification-induced correlations. Following the notation of [50], there is a "GG" correlation, in which both the spectro and photo galaxy samples are lensed by mass fluctuations in the foreground of both. Then there is a "GI" effect, in which the mass associated with the fluctuations in the foreground sample induces magnification bias on the background sample.

The GI correlation is as follows: Let the spectroscopic bin z_s span a range $\Delta\chi_s$ in comoving radial distance. The matter fluctuations m_s induce a lensing convergence on the photo-z bin at z_p of

$$\kappa_p = \frac{3\omega_m}{2}\Delta\chi_s \frac{\chi_p - \chi_s}{\chi_p} m_s, \tag{35}$$

where $\omega_m = \Omega_m h^2 = 0.127$ is the comoving matter density, χ_p and χ_s are the comoving angular diameter distances to z_s and z_p, and we have assumed a flat Universe.

The lensing magnification will induce apparent density fluctuations in the background sample as

$$g_p^{\rm lens} = q_p \kappa_p, \tag{36}$$

where q_p is a magnification bias factor for the galaxies in the photo-z bin. For instance if the selection criteria for the bin were a simple flux limit, and the intrinsic flux distribution were a power law $dn/df \propto f^{-s}$, then we would have $q_p = 2s - 2$. In general q_p will be of order unity.

The foreground galaxy distribution g_s has a correlation coefficient r_s with the mass m_s, hence a covariance between populations results:

$$\begin{aligned} \langle g_s g_p \rangle_{GI} &= \frac{3\omega_m}{2}\Delta\chi_s \frac{\chi_p - \chi_s}{\chi_p} q_p b_s r_s \langle m_s^2 \rangle & (37) \\ &= \frac{3\omega_m}{2}\Delta\chi_s \frac{\chi_p - \chi_s}{\chi_p} q_p \frac{r_s}{b_s} \langle g_s^2 \rangle & (38) \end{aligned}$$

(we have ignored shot noise in the galaxy auto-correlation). This lensing contamination will have to be subtracted from $\langle g_s g_p \rangle$ in order to extract the information on contamination c_{sp}. Even if all the cosmological factors are well determined, the magnification coefficient q_p will have to be empirically estimated. Finite accuracy in this estimate will increase δc_{sp}.

The GG correlation scales as follows: let κ_s be the convergence induced on the foreground (spectroscopic) sample by mass at $z < z_s$. This produces density fluctuations $g_s = q_s \kappa_s$. This mass induces convergence $\kappa_p \approx \kappa_s r(\chi_s, \chi_p)$ on the background (photo-z) source population, where r is an integral involving the distributions of foreground mass which must satisfy $r \geq 1$. Not concerning ourselves with details, we take $r = 1$. The induced angular correlation will be

$$\langle g_s g_p \rangle_{GG} = q_s q_p r \langle \kappa_s^2 \rangle. \tag{39}$$

Typical RMS values κ_s are 0.01–0.02 at cosmological distances. The GG lensing correlation must be removed from the signal to retrieve the contamination fraction, and again even if there is no shot noise and all distances and lensing amplitudes are known perfectly, the values of q_s and q_p will only be known to finite precision.

C. Estimate of systematic errors

Summing the GG, GI, and intrinsic contributions, the cross-correlation between spectroscopic and photometric samples is

$$\begin{aligned} \langle g_s g_p \rangle &= \left\{ \frac{3\omega_m}{2}\Delta\chi_s \frac{\chi_p - \chi_s}{\chi_p} q_p \frac{r_s}{b_s} + c_{sp} \frac{b_{sp} r_{sp}}{b_s} \right\} \langle g_s^2 \rangle + q_s q_p r \langle \kappa_s^2 \rangle & (40) \\ \Rightarrow c_{sp} &= \frac{1}{b_{sp} r_{sp}} \left[\frac{b_s \langle g_s g_p \rangle}{\langle g_s^2 \rangle} - \frac{b_s q_s q_p \langle \kappa_s^2 \rangle}{\langle g_s^2 \rangle} - \frac{3\omega_m}{2}\Delta\chi_s \frac{\chi_p - \chi_s}{\chi_p} r_s q_p \right]. & (41) \end{aligned}$$

All of the right-hand quantities are potentially well measured from the survey data itself or from other cosmological probes, except the outlier covariance factor $b_{sp} r_{sp}$ and the magnification coefficients q_p and q_p. Uncertainties in the *a priori* assumed values of these factors will propagate into the contamination coefficient as

$$(\delta c_{sp})^2 \approx \left[\delta(b_{sp} r_{sp})\right]^2 c_{sp}^2 + \delta q_p^2 \left(\frac{3\omega_m}{2}\Delta\chi_s \frac{\chi_p - \chi_s}{\chi_p} \right)^2 + \left(\delta q_p^2 + \delta q_s^2\right) \left(\frac{q \langle \kappa_s^2 \rangle}{\langle g_s^2 \rangle} \right)^2 \tag{42}$$

Here we assume $b \approx r \approx 1$, $q_s \approx q_p$.

Earlier we showed that contamination into an outlier "island" should be known to $\delta c_{sp} \leq 0.0015$ to avoid significant parameter bias. Can such a small contamination be measured using the cross-correlation technique?

- If the nominal outlier rate is $c_{sp} \approx 10\delta c_{sp} \approx 0.015$, then we require a prior estimate of outlier bias/covariance accurate to $\delta(b_{sp} r_{sp}) < 0.1$. Little will be known about the outlier population besides its luminosity range, and the outliers may tend to be active galaxies or those with unusual spectra whose clustering properties might be deviant as well. We would consider a 10% prior knowledge on outlier bias to be optimistic but perhaps attainable.

- For the second (GI) term, if we take the distance factors to be ≈ 1, and the outlier population to span a range $\Delta\chi_s \approx 0.3$, then the magnification bias coefficient must be known to an accuracy of $\delta q_p \lesssim 0.025$. This accuracy in q_p will be challenging to achieve. If the galaxy selection is by simple magnitude cut, then the slope of the counts yields q_p and potentially could be measured to high precision. Weak lensing samples typically have more complex cuts and weightings, however, than a simple flux cutoff. Surface brightness, photo-z accuracy, and ellipticity errors are involved, making estimation of q_p more difficult.

- The third (GG) term places constraints on δq_p and δq_s that will generally be weaker than those from the GI term.

If the cross-correlation technique is to determine outlier contamination fractions to an accuracy that renders them harmless, then we will need to know the $b_{sp} r_{sp}$ product of the outlier population to 0.1 or better, and also must know the magnification-bias coefficients q of our populations to 2% accuracy. This is true regardless of sample size, and these tolerances will scale as $f_{\rm sky}^{-1/2}$. The demands on $\delta(b_{sp} r_{sp})$ also becomes more stringent linearly with the photo-z outlier rate.

We have not considered the possibility of extraneous angular correlations induced by dust correlated with the foreground galaxy sample, or by dust in front of both samples (B. Menard, private communication). This signal will be present to some degree, though may perhaps be diagnosed with color information.

In summary, while we have shown here that the cross-correlation technique proposed by [51] is sensitive to catastrophic redshift errors, we found that prospects of measuring these errors (that is, the contamination coefficients c_{sp}) will be difficult using this technique alone.

VI. DISCUSSION AND CONCLUSION

In this paper we have considered the effects of the previously ignored catastrophic redshift errors — cases when the photometric redshift is grossly misestimated, i.e. when $|z_p - z_s| \sim O(1)$, and are represented by arbitrary "islands" in the $z_p - z_s$ plane. We developed a formalism, captured by Eqs. (13), that treats these islands as small "leakages" (or "contaminations") and directly estimates their effect on bias in cosmological parameters. We then inverted the problem by estimating how many spectroscopic redshifts are required to control catastrophic errors at a level that makes them harmless for cosmology. In the process, we have proven two general-purpose theorems (in the Appendices): that the bias due to systematics always decreases or stays fixed if 1) (unbiased) prior information is added to the fiducial survey, or 2) we marginalize over one or more dimensions of the parameter space.

We found that, at face value, of order million redshifts are required in order not to bias the dark energy parameter measurements (that is, in order to lead to $\Delta\chi^2 \lesssim 1$ in the $w_0 - w_a$ plane). However, the requirement becomes significantly (30 times) less stringent if we restrict the survey to redshift $z < 2.5$; in that case, $N_{\rm spec}$ is only of order a few tens of thousands. Essentially, leakage of galaxies from lower redshift to $z > 2.5$ is damaging since there are few galaxies at such high redshift and the relative bias in galaxy number is large. Therefore, using only galaxies with $z \lesssim 2.5$ helps dramatically by lowering the required $N_{\rm spec}$ while degrading the dark energy (figure-of-merit) constraints by mere $\sim 20\%$.

We have studied two approaches for a spectroscopic survey: the untargeted one where the number of spectra at each redshift is proportional to the number of photometric galaxies (§IV A) and the targeted one where the number of spectra is optimized to be minimal for a given degradation in cosmological parameters (§IV B). For the case where galaxies with $z_p > 2.5$ are dropped, the targeted survey gave only a modestly ($\sim 40\%$) smaller required $N_{\rm spec}$.

We do not imply that these $N_{\rm spec}$ requirements to apply to all proposed surveys to high accuracy, although the $O(10^{-3})$ required knowledge on catastrophe rates is robust. The calculation should be repeated with the fiducial photo-z outlier distribution, survey characteristics, and cosmological parameters of interest to a particular experiment.

Our work demonstrates for example that efforts to reduce the "island" of catastrophic mis-assignment from $z \approx 0$ to $z \approx 3$, such as magnitude priors, could greatly reduce the required $N_{\rm spec}$. Since $N_{\rm spec} \propto f_{\rm sky}\bar{c}$ (with $\bar{c}$ being the mean rate of catastrophic contamination), it is clear that a photo-z survey with improved S/N and wavelength coverage to reduce the total catastrophe rate will also require lower $N_{\rm spec}$ to calibrate these rates.

Another practical implication of these results is that the spectroscopic redshift surveys must be of very high completeness—99.9% if there is a possibility that all failures could be in an outlier island, but less if some fraction of the failures are known to be in the core of the error distribution.

If an outlier island is known to exist at a particular (z_s, z_p) location, it may be possible to include the contamination c_{sp} as a free parameter in the data analysis and marginalize over its value. It is possible that self-calibration may reduce the bias in cosmological parameters. It is likely infeasible, however, to leave the c_{sp} values over the full (z_s, z_p) plane as free parameters. We leave self-calibration of outlier rates for future work.

We have also studied whether the technique proposed by Newman [51], which correlates a photometric sample with a spectroscopic one, can be used to measure, and thus correct for, catastrophic redshift errors. The advantage of this approach is that the spectroscopic survey need not be a representative sampling of the photometric catalog. While we found that the cross-correlation technique is sensitive to catastrophic errors (specifically, the contamination coefficients c_{sp}), the contamination coefficient is degenerate with the value $b_{sp}r_{sp}$ of the bias and stochasticity of the outlier population. Furthermore there is a correlation induced by lensing magnification bias that spoofs the contamination signal. It will therefore be difficult to use the cross-correlation technique to constrain outlier rates to the requisite accuracy.

Overall, we are very optimistic that the catastrophic redshift errors can be controlled to the desired accuracy. We have identified a simple strategy that requires only of order 30,000 spectra out to $z \simeq 2.5$ for the calibration to be successful for a SNAP-type survey. Incidentally, this number of spectra required for the catastrophic errors is of the same order of magnitude as that required for the non-catastrophic, "core" errors [11, 12, 14]. Total spectroscopic requirements of a survey will be based on the greater of requirements of these two error regimes.

VII. ACKNOWLEDGEMENTS

GB is supported by grant AST-0607667 from the NSF, DOE grant DE-FG02-95ER40893, and NASA grant BEFS-04-0014-0018. DH is supported by the DOE OJI grant under contract DE-FG02-95ER40899, NSF under contract AST-0807564, and NASA under contract NNX09AC89G.

APPENDIX A: EFFECT OF UNBIASED PRIORS ON BIAS SIGNIFICANCE

Will a bias get worse or better (more or less significant) when additional unbiased prior information is added to the likelihood? It is intuitive that biases ΔP should decrease when unbiased information is added. However $F \rightarrow F+G$ for some new non-negative-definite prior Fisher matrix G, meaning that the statistical errors also shrink. So which effect wins out? We prove here that addition of unbiased prior information *cannot increase* the significance of parameter bias.

The proof is straightforward: Eq. (9) gives the significance $\Delta\chi^2$ of a bias in terms of the original positive-definite Fisher matrix F and the vector V. If G is a non-negative-definite prior, then the change significance of the bias is

$$\Delta\chi^2(\text{with prior}) - \Delta\chi^2(\text{without prior}) = V^T(F+G)^{-1}V - V^TF^{-1}V \tag{A1}$$

This quantity cannot be positive. If it were, then there would some $0 < \lambda_0 < 1$ and a positive-definite matrix $H = F + \lambda_0 G$ such that

$$0 < \frac{\partial}{\partial\lambda}\left[V^T(F+\lambda G)^{-1}V\right]_{\lambda_0} = -(H^{-1}V)^T G H^{-1} V. \tag{A2}$$

If G is non-negative definite, this situation cannot occur. We hence conclude that the $\Delta\chi^2$ of some bias is always reduced (or stays the same) by addition of an unbiased prior.

APPENDIX B: EFFECT OF MARGINALIZATION ON BIAS SIGNIFICANCE

Second we can ask: If we calculate the significance of a bias induced over a parameter space, then marginalize away parameter vector B to leave parameter vector A, how might the significance differ in the smaller space? We show that in the Gaussian limit, *marginalization always reduces (or leaves unchanged) the $\Delta\chi^2$ assigned to the bias, although the $\Delta\chi^2$ per DOF may increase.* To see this: first we note that marginalization over B does not change the biases in the parameters A if the distribution is Gaussian. So the bias in A is simply a projection matrix P_A times ΔP: $\Delta P_A = P_A F^{-1} V$. The $\Delta\chi^2$ after marginalization down to the A space is determined by the marginalized Fisher matrix, $F' = [(F^{-1})_{AA}]^{-1}$. So we have

$$\begin{aligned}(\Delta\chi^2)_A &= V^T F^{-1} P_A^T F' P_A F^{-1} V \\ &= V^T F^{-1} V - (P_B V)^T (F_{BB})^{-1} (P_B V) \qquad \text{(B1)} \\ &= \Delta\chi^2 - (P_B V)^T (F_{BB})^{-1} (P_B V). \qquad \text{(B2)}\end{aligned}$$

The equivalence in (B1) can be derived from manipulation of the common expression for the inverse of a matrix decomposed into an 2×2 array of submatrices. Because F_{BB} and its inverse must be non-negative-definite, the last term is negative, so we are assured that $(\Delta\chi^2)_A \leq \Delta\chi^2$. Equality is, however, easily obtained, for example if there is no bias in the B parameters. We thus know that $\Delta\chi^2/N_{\rm DOF}$ can potentially increase.

[1] M. Bartelmann and P. Schneider, Phys. Rept. **340**, 291 (2001), astro-ph/9912508.
[2] D. Huterer, Phys. Rev. **D65**, 063001 (2002), astro-ph/0106399.
[3] D. Munshi, P. Valageas, L. Van Waerbeke, and A. Heavens, Phys. Rept. **462**, 67 (2008), astro-ph/0612667.
[4] H. Hoekstra and B. Jain (2008), arXiv:0805.0139.
[5] A. Albrecht et al. (2009), arXiv:0901.0721.
[6] H. Oyaizu, M. Lima, C. E. Cunha, H. Lin, and J. Frieman (2007), arXiv:0711.0962.
[7] M. Banerji, F. B. Abdalla, O. Lahav, and H. Lin (2007), arXiv:0711.1059.
[8] M. Lima et al., Mon. Not. Roy. Astron. Soc. **390**, 118 (2008), arXiv:0801.3822.
[9] C. E. Cunha, M. Lima, H. Oyaizu, J. Frieman, and H. Lin (2008), arXiv:0810.2991.
[10] F. B. Abdalla, M. Banerji, O. Lahav, and V. Rashkov (2008), arXiv:0812.3831.
[11] Z.-M. Ma, W. Hu, and D. Huterer, Astrophys. J. **636**, 21 (2005), astro-ph/0506614.
[12] D. Huterer, M. Takada, G. Bernstein, and B. Jain, Mon. Not. Roy. Astron. Soc. **366**, 101 (2006), astro-ph/0506030.
[13] T. D. Kitching, A. N. Taylor, and A. F. Heavens (2008), arXiv:801.3270.
[14] Z. Ma and G. Bernstein, Astrophys. J. **682**, 39 (2008), arXiv:0712.1562.
[15] D. Huterer and M. Takada, Astropart. Phys. **23**, 369 (2005), astro-ph/0412142.
[16] K. Heitmann, P. M. Ricker, M. S. Warren, and S. Habib, Astrophys. J. Suppl. **160**, 28 (2005), astro-ph/0411795.
[17] K. Heitmann, M. White, C. Wagner, S. Habib, and D. Higdon (2008), arXiv:0812.1052.
[18] A. R. Zentner, D. H. Rudd, and W. Hu, Phys. Rev. **D77**, 043507 (2008), arXiv:0709.4029.
[19] M. Takada and B. Jain (2008), arXiv:0810.4170.
[20] M. J. White, Astropart. Phys. **22**, 211 (2004), astro-ph/0405593.
[21] B. Hagan, C.-P. Ma, and A. V. Kravtsov, Astrophys. J. **633**, 537 (2005), astro-ph/0504557.
[22] H. Zhan and L. Knox, Astrophys. J. **616**, L75 (2004), astro-ph/0409198.
[23] D. Huterer and M. J. White, Phys. Rev. **D72**, 043002 (2005), astro-ph/0501451.
[24] C. Shapiro and A. Cooray, JCAP **0603**, 007 (2006), astro-ph/0601226.
[25] C. Shapiro (2008), arXiv:0812.0769.
[26] D. H. Rudd, A. R. Zentner, and A. V. Kravtsov (2007), astro-ph/0703741.
[27] T. D. Kitching, A. Amara, F. B. Abdalla, B. Joachimi, and A. Refregier (2008), arXiv:0812.1966.
[28] B. Joachimi and P. Schneider (2008), arXiv:0804.2292.
[29] G. M. Bernstein (2008), arXiv:0808.3400.
[30] J. Guzik and G. Bernstein, Phys. Rev. **D72**, 043503 (2005), astro-ph/0507546.
[31] C. Heymans et al., Mon. Not. Roy. Astron. Soc. **368**, 1323 (2006), astro-ph/0506112.
[32] R. Massey et al., Mon. Not. Roy. Astron. Soc. **376**, 13 (2007), astro-ph/0608643.
[33] S. Paulin-Henriksson, A. Amara, L. Voigt, A. Refregier, and S. L. Bridle (2007), arXiv:0711.4886.
[34] A. Amara and A. Refregier (2007), arXiv:0710.5171.
[35] H. F. Stabenau, B. Jain, G. Bernstein, and M. Lampton (2007), arXiv:0710.3355.
[36] M. Jarvis and B. Jain (2004), astro-ph/0412234.
[37] M. Takada and S. Bridle, New J. Phys. **9**, 446 (2007), arXiv:705.0163.
[38] Z. Ma, G. Bernstein, A. Weinstein, and M. Sholl (2008), arXiv:809.2954.
[39] A. Heavens, A. Refregier, and C. Heymans, Mon. Not. Roy. Astron. Soc. **319**, 649 (2000), astro-ph/0005269.
[40] R. A. C. Croft and C. A. Metzler (2000), astro-ph/0005384.
[41] R. G. Crittenden, P. Natarajan, U.-L. Pen, and T. Theuns, Astrophys. J. **559**, 552 (2001), astro-ph/0009052.
[42] J. Mackey, M. J. White, and M. Kamionkowski, Mon. Not. Roy. Astron. Soc. **332**, 788 (2002), astro-ph/0106364.
[43] Y. P. Jing, Mon. Not. Roy. Astron. Soc. **335**, L89 (2002), astro-ph/0206098.
[44] C. Heymans and A. Heavens, Mon. Not. Roy. Astron. Soc. **339**, 711 (2003), astro-ph/0208220.
[45] L. J. King and P. Schneider, Astron. Astrophys. **398**, 23 (2003), astro-ph/0209474.
[46] C. M. Hirata et al. (SDSS), Mon. Not. Roy. Astron. Soc. **353**, 529 (2004), astro-ph/0403255.

[47] R. Mandelbaum, C. M. Hirata, M. Ishak, U. Seljak, and J. Brinkmann, Mon. Not. Roy. Astron. Soc. **367**, 611 (2006), astro-ph/0509026.
[48] S. Bridle and F. B. Abdalla, Astrophys. J. **655**, L1 (2007), astro-ph/0608002.
[49] S. Bridle and L. King, New J. Phys. **9**, 444 (2007), arXiv:705.0166.
[50] C. M. Hirata and U. Seljak, Phys. Rev. **D70**, 063526 (2004), astro-ph/0406275.
[51] J. A. Newman (2008), arXiv:0805.1409.
[52] L. Knox, R. Scoccimarro, and S. Dodelson, Phys. Rev. Lett. **81**, 2004 (1998), astro-ph/9805012.
[53] D. Huterer and M. S. Turner, Phys. Rev. **D64**, 123527 (2001), astro-ph/0012510.
[54] M. Tegmark, A. Taylor, and A. Heavens, Astrophys. J. **480**, 22 (1997), astro-ph/9603021.
[55] G. Aldering et al. (SNAP) (2004), astro-ph/0405232.
[56] A. Albrecht, G. Bernstein, R. Cahn, W. L. Freedman, J. Hewitt, W. Hu, J. Huth, M. Kamionkowski, E. W. Kolb, L. Knox, et al. (2006), astro-ph/0609591.
[57] E. Komatsu et al. (WMAP) (2008), arXiv:0803.0547.
[58] S. Jouvel et al. (2009), arXiv:0902.0625.
[59] T. Dahlen, B. Mobasher, S. Jouvel, J.-P. Kneib, O. Ilbert, S. Arnouts, G. Bernstein, and J. Rhodes, Astrophys. J. **136**, 1361 (2008), 0710.5532.
[60] M. Schneider, L. Knox, H. Zhan, and A. Connolly, Astrophys. J. **651**, 14 (2006), astro-ph/0606098.